



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/662,108	09/12/2003	Geoffrey M. Brown	52224/297580	9099

23370 7590 05/14/2008
JOHN S. PRATT, ESQ
KILPATRICK STOCKTON, LLP
1100 PEACHTREE STREET
ATLANTA, GA 30309

EXAMINER

TURNER, ASHLEY D

ART UNIT	PAPER NUMBER
----------	--------------

2154

MAIL DATE	DELIVERY MODE
-----------	---------------

05/14/2008

PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary

Application No.

10/662,108

Applicant(s)

BROWN, GEOFFREY M.

Examiner

ASHLEY D. TURNER

Art Unit

2154

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period **will** apply and **will** expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply **will**, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 12 September 2003.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-50 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-50 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. _____.
3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☒ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) ☒ Information Disclosure Statement(s) (PTO/SB/08)
Paper No(s)/Mail Date 12/11/2003.
- 4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date: _____.
- 5) ☐ Notice of Informal Patent Application
- 6) ☐ Other: _____.

DETAILED ACTION

Claim Rejections - 35 USC § 102

1. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

Claim 1-50 are rejected under 35 U.S.C. 102 (b) as being anticipated by Lloyd et al hereinafter Lloyd (US 2003/0039212 A1).

Referring to claim 1 Lloyd discloses a method of clustering a plurality of network destinations having network addresses being portioned into groups of network addresses according to an initial grouping, comprising : identifying a plurality of seedpoints from among the network destinations, each seedpoint being an active one of the destinations associated with at least one of the groups of network addresses; topologically clustering the seedpoints into groups of topologically similar seedpoints; performing a measurement from a predetermined location to a seedpoint within each group of seedpoints; clustering the seedpoints into clusters based on the measurements, the clusters being the number of clusters and the similarity among the measurements for the seedpoints within each cluster and generalizing the clusters based on information identifying the network addresses with corresponding seedpoints to which the network addresses are deemed close. (Pg. 1 Paragraphs[0007] Assessment and Optimization of network traffic is also subjective: performance characteristics toward a given destination vary based on the location from

Art Unit: 2154

which measurements are made. Each locale must identify their own local performance characteristics. [0008] As the number of valid destination addresses increase, the feasibility of measuring performance characteristics toward all destinations in a persistent and ongoing manner decreases to nil. [0009] Consolidating the set of all individual destination addresses into sets, or ranges of addresses significantly reduces the number of tests that must be made in a given time interval. However, in many cases, even this optimization is insufficient: the current Internet routing tables include over 100,000 different "routes", or address range identifiers. [0010] The performance aspects of network paths to any given address will often have similar characteristics of the network paths to a nearby address. However, the task of identifying address blocks which are acceptably similar in performance characteristic is subtle: network address ranges are typically not directly representative of the underlying network topology. [0011] In most locations, the set of actively used address ranges is a relatively small percentage of the total address ranges available. This set of "active" routes is dynamic, and while technology does exist that can identify the set of active flows, in general these systems are constrained by both significant delays in reporting activity, as well as lack of integration into path selection systems.)

Regarding claim 2

Referring to claim 2 Lloyd discloses all the limitations of claim 2 which is described above. Lloyd also discloses wherein the seedpoints are identified based on a predetermined desired granularity. (Pg, 7 paragraph [0162] In the implementation described in this document, low resolution in the metric may be considered acceptable,

Art Unit: 2154

since this tends to diminish the likelihood of reacting to small changes in performance. The goal of some embodiments is to find several routes, none of which are appreciably sub-optimal, and any one of which can be pulled out of service rapidly (with the others taking over) when sudden performance degradations such as packet drop storms are observed. In those embodiments, paths whose measurements vary by insignificant amounts are treated as equal; in some implementations, this can be achieved with a composite metric of low granularity.)

Regarding claim 3

Referring to claim 3 Lloyd discloses all the limitations of claim 3 which is described above. Lloyd also discloses wherein the predetermined desired granularity is expressed as a number of most significant bits of the destination addresses. (Pg. Paragraph 1[0009]

Consolidating the set of all individual destination addresses into sets, or ranges of addresses significantly reduces the number of tests that must be made in a given time interval. However, in many cases, even this optimization is insufficient: the current Internet routing tables include over 100,000 different "routes", or address range identifiers.)

Regarding claim 4

Referring to claim 4 Lloyd discloses all the limitations of claim 4 which is described above. Lloyd also discloses wherein the destinations addresses are 32 bits in length, and predetermined number of most significant bits is 24. (Pg. 8 [0193] In one such implementation, each PoP will be equipped with an implementation of BGP. These BGP

Art Unit: 2154

sessions should communicate with each other; in some implementations, all PoPs can peer in a full mesh of IBGP peerings. In other implementations where the number of PoPs is large so that a full mesh becomes impractical, BGP Route Reflection can be implemented. In such an infrastructure, each PoP can advertise routes to all EC's that it has monitored, giving an address within that PoP as the next hop. Most attributes in that BGP message can be left at default values; in some implementations, only the 32-bit Local Preference field is modified as to communicate the measurements described above for the closeness of the PoP to the EC.) [0309] In some embodiments of the invention, the measurement packets are of varying sizes, including 64, 256, 512, 1024, 1500 bytes.

Regarding claim 5

Referring to claim 5 Lloyd discloses all the limitations of claim 5 which is described above. Lloyd also discloses wherein the density of the seedpoints in different groups is varied based upon traffic statistics.(Pg. 2 [0044] Monitoring Exit Traffic [0045] In this embodiment of the invention, a passive monitoring and reporting device is needed, either in the enterprise DMZ or in a collocation facility (CoLo) rack; let's call this device the Collector. The job of the collector is to maintain simple state on selected flows emerging from the CP, and report statistics on them. In some embodiments of this invention, there is an Ethernet link to instrument, in which case packet splitters can be used; to such a device a computer that runs the Collector software can be attached. The software needs to receive signals about addresses (and possibly ports) for which to monitor flows.

Art Unit: 2154

Regarding claim 6

Referring to claim 6 Lloyd discloses all the limitations of claim 6 which is described above. Lloyd also discloses wherein the seedpoints are identified based on information concerning the destinations of data traffic in the network.(Pg. 2 [0048] In some embodiments of this invention, the device can select flows on the basis of input other than performance, such as user input, or the cost of the link. That is, the user can specify the flows that need special treatment. Also, interface statistics can be obtained from network; using this input, quantities such as available bandwidth, cost of usage of the links, or link bandwidth can be computed. Those skilled in the art can identify other metrics based on which flow determination and treatment can be based. Flows can be diverted when a given threshold in such metrics (such as cost or available bandwidth) is detected. (Pg. 3 [0057] In other embodiments of this invention, flows are identified by matching criteria defined on any portion of an HTTP request. In one implementation of such an embodiment, the well-known technique of identifying individuals or groups using cookies can be used.

Regarding claim 7

Referring to claim 7 Lloyd discloses all the limitations of claim 7 which is described above. Lloyd also discloses wherein the seedpoints are identified by sending a message to at least one address in each of a complete set of address regions spanning the destination addresses, each region being defined by a corresponding unique pattern of most significant address bits. (Pg. 8 [0193] In one such implementation, each PoP will be

Art Unit: 2154

equipped with an implementation of BGP. These BGP sessions should communicate with each other; in some implementations, all PoPs can peer in a full mesh of IBGP peerings. In other implementations where the number of PoPs is large so that a full mesh becomes impractical, BGP Route Reflection can be implemented. In such an infrastructure, each PoP can advertise routes to all EC's that it has monitored, giving an address within that PoP as the next hop. Most attributes in that BGP message can be left at default values; in some implementations, only the 32-bit Local Preference field is modified as to communicate the measurements described above for the closeness of the PoP to the EC. Those skilled in the art can identify other fields that can be used for this purpose (e.g., weight, MED, etc.) In implementations that use Local Preference, desirable routes should have higher values; in some embodiments, a scheme based on an ideal number of points from which "demerits" are subtracted, for example when packets are observed to drop, or delay crosses a series of thresholds, can be implemented. (Pg. 13 [0259] In some embodiments of the invention, this control information will include notification of measurement devices that have joined or left the network.)

Regarding claim 8

Referring to claim 8 Lloyd discloses all the limitations of claim 8 which is described above. Lloyd also discloses wherein the one address in of the address regions is one of a set predetermined addresses within each region to which messages are conditionally sent to identify seedpoints. (Pg.8 [0193] In one such implementation, each PoP will be equipped with an implementation of BGP. These BGP sessions should communicate with

Art Unit: 2154

each other; in some implementations, all PoPs can peer in a full mesh of IBGP peerings. In other implementations where the number of PoPs is large so that a full mesh becomes impractical, BGP Route Reflection can be implemented. In such an infrastructure, each PoP can advertise routes to all EC's that it has monitored, giving an address within that PoP as the next hop. Most attributes in that BGP message can be left at default values; in some implementations, only the 32-bit Local Preference field is modified as to communicate the measurements described above for the closeness of the PoP to the EC. Those skilled in the art can identify other fields that can be used for this purpose (e.g., weight, MED, etc.) In implementations that use Local Preference, desirable routes should have higher values; in some embodiments, a scheme based on an ideal number of points from which "demerits" are subtracted, for example when packets are observed to drop, or delay crosses a series of thresholds, can be implemented.

Regarding claim 9

Referring to claim 9 Lloyd discloses all the limitations of claim 9 which is described above. Lloyd discloses wherein the seedpoints are included in autonomous systems (Pg. 2 [0042] In the context of one specific embodiment of this invention, we describe the hardware needed to collect measurements and bend traffic appropriately. Whether the mechanisms are implemented at a DMZ or Colo, the existing network infrastructure in which these mechanisms are being implemented can either be an autonomous system, or an autonomous sub-system.), further comprising selecting a representative for each

Art Unit: 2154

cluster of seedpoints, the selecting of representative including determining whether the representative has the same penultimate hop along a path to the autonomous system as do the seedpoints of the cluster. (Pg.8 [0193] In one such implementation, each PoP will be equipped with an implementation of BGP. These BGP sessions should communicate with each other; in some implementations, all PoPs can peer in a full mesh of IBGP peerings. In other implementations where the number of PoPs is large so that a full mesh becomes impractical, BGP Route Reflection can be implemented. In such an infrastructure, each PoP can advertise routes to all EC's that it has monitored, giving an address within that PoP as the next hop. Most attributes in that BGP message can be left at default values; in some implementations, only the 32-bit Local Preference field is modified as to communicate the measurements described above for the closeness of the PoP to the EC. Those skilled in the art can identify other fields that can be used for this purpose (e.g., weight, MED, etc.) In implementations that use Local Preference, desirable routes should have higher values; in some embodiments, a scheme based on an ideal number of points from which "demerits" are subtracted, for example when packets are observed to drop, or delay crosses a series of thresholds, can be implemented.

Regarding claim 10

Referring to claim 10 Lloyd discloses all the limitations of claim 10 which is described above. Lloyd discloses wherein identifying seedpoints includes rejecting those seedpoints for which there is no available representative in the same autonomous system. (Pg. 2 [0042] In the context of one specific embodiment of this invention, we describe the

Art Unit: 2154

hardware needed to collect measurements and bend traffic appropriately. Whether the mechanisms are implemented at a DMZ or Colo, the existing network infrastructure in which these mechanisms are being implemented can either be an autonomous system, or an autonomous sub-system.)(Pg. 1 [0016] A method of classifying network addresses into appropriate Groups, based on evaluation of a selection of performance, policy, topology, and related criteria. Having defined said Groups, instantiate them in a server or other networking device, such that performance scores for observed traffic can be consolidated within the Group definition.)

Regarding claim 11

Referring to claim 11 Lloyd discloses all the limitations of f claim 10 which is described above. Lloyd also discloses wherein topologically clustering comprises performing traceroute operations to the seedpoints and analyzing the resulting reported routes.(Pg. 4 [0088] In the DMZ case, some embodiments of this invention involve having a process that adds static routes to the router (in addition to those static routes advertised to divert the flows to the Bender, described above). These processes may involve the addition of several specific statics, with next hops that spread across the GRE tunnels leading to various PoPs. In some embodiments of the invention, this indirection is outside the GRE tunnels themselves, hence preserving the property that a single flow will all go to the same next PoP, even if load sharing across tunnels per flow is in use. (Pg. 6 [0144] 2. Traceroute towards the Surfer. [0145] 3. Coercing the Surfer to communicate with a PoP as part of the conversation, for example by embedding a particular GIF as part of their

Art Unit: 2154

Web page. (Telnet or other protocols could also be used.) [0146] 4. Asking the Surfer to run some code that initiates connections to candidate PoPs. (The motivation for the Surfer would be the improved performance if they were to do this. Note also that only one or a few Surfers within an EC need to accept the code to get a measurement for the whole EC.)

Regarding claim 12

Referring to claim 12 Lloyd discloses all the limitations of claim 12 which is described above. Lloyd also discloses wherein the measurement performed from the predetermined location to each of the seedpoints is one of multiple measurements performed from the predetermined location to each of the seedpoints. (Pg. 1 [0007] Assessment and Optimization of network traffic is also subjective: performance characteristics toward a given destination vary based on the location from which measurements are made. Each locale must identify their own local performance characteristics.)

Regarding claim 13

Referring to claim 13 Lloyd discloses all the limitations of claim 13 which is described above. Lloyd also discloses wherein the predetermined location is one of multiple predetermined locations from which measurements to the seedpoints are performed. (Pg. 1[0007] Assessment and Optimization of network traffic is also subjective: performance characteristics toward a given destination vary based on the location from which

Art Unit: 2154

measurements are made. Each locale must identify their own local performance characteristics.)

Regarding claim 14

Referring to claim 14 Lloyd discloses all the limitations of claim 14 which is described above. Lloyd also discloses wherein performing each measurement comprises sending a time-to-live limited probe message to a candidate seedpoint. (Pg. 4 [0076] Some embodiments implement Option 7. A simple web interface is offered as part of the Bender. Additional logic will take care of performing post operations on the web transactions. A secure web form interface can also be offered as part of other Options for flow qualification; this allows the Content Provider to either experiment with the service, or intervene to control it once it is running. In a simplified embodiment of this invention, the HTML post interface supports one operation, namely to deem one IP address a qualified address for a specified period of time. This period of time is specified by the Content Provider, which enters for this purpose a TTL (Time To Live) for the qualification; it may be desirable to keep the TTL value small; it may also be desirable to provide a default, possibly at around 5 minutes.)

Regarding claim 15

Referring to claim 15 Lloyd discloses all the limitations of claim 15 which is described above. Lloyd also discloses wherein performing each measurement comprises sending an

Art Unit: 2154

echo request message to a candidate seedpoint. (Pg. 8 [0193] In one such implementation, each PoP will be equipped with an implementation of BGP. These BGP sessions should communicate with each other; in some implementations, all PoPs can peer in a full mesh of IBGP peerings. In other implementations where the number of PoPs is large so that a full mesh becomes impractical, BGP Route Reflection can be implemented. In such an infrastructure, each PoP can advertise routes to all EC's that it has monitored, giving an address within that PoP as the next hop. Most attributes in that BGP message can be left at default values; in some implementations, only the 32-bit Local Preference field is modified as to communicate the measurements described above for the closeness of the PoP to the EC. Those skilled in the art can identify other fields that can be used for this purpose (e.g., weight, MED, etc.) In implementations that use Local Preference, desirable routes should have higher values; in some embodiments, a scheme based on an ideal number of points from which "demerits" are subtracted, for example when packets are observed to drop, or delay crosses a series of thresholds, can be implemented.

Regarding claim 16

Referring to claim 16 Lloyd discloses all the limitations of claim 16 which is described above. Lloyd also discloses wherein the measurements are temporal measurements. (Pg. 1[0007] Assessment and Optimization of network traffic is also subjective: performance characteristics toward a given destination vary based on the location from which measurements are made. Each locale must identify their own local performance characteristics.)

Regarding claim 17

Referring to claim 17 Lloyd discloses all the limitations of claim 17 which is described above. Lloyd also discloses wherein clustering the seedpoints includes ordering the seedpoints according to the measurement. (Pg. 1 [0016] A method of classifying network addresses into appropriate Groups, based on evaluation of a selection of performance, policy, topology, and related criteria. Having defined said Groups, instantiate them in a server or other networking device, such that performance scores for observed traffic can be consolidated within the Group definition.)

Regarding claim 18

Referring to claim 18 Lloyd discloses all the limitations of claim 18 which is described above. Lloyd also discloses each cluster of ordered seedpoints, identifying at least one of the seedpoints whose measurement satisfies a predetermined criterion. (Pg. 7 [0154] In some embodiments, measurements can be boiled down into a single statistic. In some cases, it may be important to rank the measurements by their expected quality. In one example implementation, the quality order is: [0155] 1. Direct ping measurements to the Surfer [0156] 2. Spoofing races; that is, a technique where multiple PoPs respond to a query, all using the same address ("spoofing"). For some common network protocols (including at least DNS, HTTP and TCP), when the answer to the query is used, it is possible to tell which PoP's answer arrived first, and hence is the closest.

Regarding claim 19

Referring to claim 19 Lloyd discloses all the limitations of claim 19 which is described above. Lloyd also discloses wherein the predetermined criterion is being closest to a centroid of the measurements of the seedpoints. (Pg. 7 [0154] In some embodiments, measurements can be boiled down into a single statistic. In some cases, it may be important to rank the measurements by their expected quality. In one example implementation, the quality order is: [0155] 1. Direct ping measurements to the Surfer [0156] 2. Spoofing races; that is, a technique where multiple PoPs respond to a query, all using the same address ("spoofing"). For some common network protocols (including at least DNS, HTTP and TCP), when the answer to the query is used, it is possible to tell which PoP's answer arrived first, and hence is the closest.

Regarding claim 20

Referring to claim 20 Lloyd discloses all the limitations of claim 20 which is described above. Lloyd also discloses wherein clustering the seedpoints is performed on the basis of autonomous systems in which the seedpoints reside. (Pg.7 and 8 [0172] In this case, both end points of a cross-core links are under our control. In some embodiments, it is desirable to generate cross-core measurements that are comparable with the Egress PoP measurements described above, so that path comparisons can be made. In some embodiments of this invention, an overlay network can constitute the network

Art Unit: 2154

infrastructure that links these PoPs to each other. In the following implementation, we assume the use of GRE tunnels across POPs. In some embodiments, this overlay network can be set up as a single Autonomous System. In this implementation, the overlay network is set up to be single BGP Autonomous system. On the core, different quantities of interest along all direct paths between all PoPs can be directly monitored. Possible measurements cover, and are not restricted to one or more of the following: [0173] 1. One-way delay [0174] 2. One-way packet drop rate [0175] 3. One-way jitter [0176] 4. Available access bandwidth. (In some embodiments, this quantity can be estimated via the output buffer fullness on the physical interface at the sending end of the inter-PoP tunnel.) [0177] 5. Complete tunnel failure)

Regarding claim 21

Referring to claim 21 Lloyd discloses all the limitations of 21 which is described above Lloyd also discloses wherein a minimum of one cluster is established per autonomous system. (Pg.7 and 8 [0172] In this case, both end points of a cross-core links are under our control. In some embodiments, it is desirable to generate cross-core measurements that are comparable with the Egress PoP measurements described above, so that path comparisons can be made. In some embodiments of this invention, an overlay network can constitute the network infrastructure that links these PoPs to each other. In the following implementation, we assume the use of GRE tunnels across POPs. In some embodiments, this overlay network can be set up as a single Autonomous System. In this implementation, the overlay network is set up to be single BGP Autonomous system. On

Art Unit: 2154

the core, different quantities of interest along all direct paths between all PoPs can be directly monitored. Possible measurements cover, and are not restricted to one or more of the following: [0173] 1. One-way delay [0174] 2. One-way packet drop rate [0175] 3. One-way jitter [0176] 4. Available access bandwidth. (In some embodiments, this quantity can be estimated via the output buffer fullness on the physical interface at the sending end of the inter-PoP tunnel.) [0177] 5. Complete tunnel failure)

Regarding claim 22

Referring to claim 22 Lloyd discloses all limitations 22 Lloyd discloses all the limitations of 22 which is described above. Lloyd also discloses wherein clustering the seedpoints is performed on the basis of traffic to the network destinations. (Pg. 2 [0044] Monitoring Exit Traffic [0045] In this embodiment of the invention, a passive monitoring and reporting device is needed, either in the enterprise DMZ or in a collocation facility (CoLo) rack; let's call this device the Collector. The job of the collector is to maintain simple state on selected flows emerging from the CP, and report statistics on them. In some embodiments of this invention, there is an Ethernet link to instrument, in which case packet splitters can be used; to such a device a computer that runs the Collector software can be attached. The software needs to receive signals about addresses (and possibly ports) for which to monitor flows.

Regarding claim 23

Art Unit: 2154

Referring to claim 23 Lloyd discloses all limitations 23 discloses all the limitations of 23 which is described above. Lloyd also discloses wherein clustering the seedpoints is performed in at least two passes, a first pass resulting in more clusters than desired, a second pass being based on a subset of the seedpoints taken from larger ones of the clusters resulting from the first pass. (Pg. 12 [0235] Relating to an embodiment of FIG. 2, suppose we have a flow at point 1 trying to reach internetwork2, and the paths have the metrics associated with them as annotated, where lower is better and metrics can be added together. Some possibilities for determining the exit point and path to the exit point include (but are not limited to): choose the exit point first and then determine path to the exit point. In this case, choose exit point 2 as the best exit point because its path to internetwork2 has metric 100 and exit point 3 has metric 200. Given that we have chosen exit point 2, we can choose either the best direct path from point 1 to exit point 2 (500) or the best indirect path (in this case, via exit point 3, 100+200). This results in a total path metric of $100+200+100=400$. Consider both the exit point's paths and the paths to the exit point in combination. In this case, choose exit point 3, because even though the path to internetwork2 from exit point 3 is 200, the path to exit point 3 from point 1 is only 100, resulting in a total path metric of $100+200=300$. [0236] Relating to one embodiment, FIG. 3a shows an internetwork1 with a network device can decide to redirect flows directly bound for internetwork 4 to another internetwork with a network device, like internetwork 2, or to an internetwork without a network device, like internetwork 3. [0237] Relating to one embodiment, FIG. 3b shows an internetwork1 can have both flow endpoints inside of it, and can have multiple network devices deployed to

Art Unit: 2154

obtain measurements, calculate metrics, select paths, and alter routing in the internetwork.)

Regarding claim 24

Referring to claim 24 Lloyd discloses all the limitations of claim 24 which is described above. Lloyd also discloses wherein clustering the seedpoints is based on geographical information about the seedpoints.(Pg. 20 and 21 A method of populating a plurality of one or more groups with a plurality of one or more network addresses, the method comprising: selecting a plurality of one or more metrics, the plurality of one or more metrics including one or more of path performance, network cost, network usage, geographical proximity, topological proximity, and statistical similarity; creating the plurality of one or more groups in one or more network devices, network devices including one or more of servers, forwarding devices, and routing devices; populating each of the plurality of one or more groups with a subset of the plurality of one or more network addresses based on at least a classification function, the classification function at least partly determined by at least one of the plurality of one or more metrics; and including one or more network addresses from one or more groups of the plurality of one or more groups in a plurality of one or more routing tables distributed across the internetwork.)

Regarding claim 25

Art Unit: 2154

Referring to claim 25 Lloyd discloses all the limitations of claim 25 is described above. Lloyd also discloses wherein clustering the seedpoints employs a clustering budget of predetermined number of clusters for each of a predetermined fraction of the total number of seedpoints. (Pg.12 [0235] Relating to an embodiment of FIG. 2, suppose we have a flow at point 1 trying to reach internetwork2, and the paths have the metrics associated with them as annotated, where lower is better and metrics can be added together. Some possibilities for determining the exit point and path to the exit point include (but are not limited to): choose the exit point first and then determine path to the exit point. In this case, choose exit point 2 as the best exit point because its path to internetwork2 has metric 100 and exit point 3 has metric 200. Given that we have chosen exit point 2, we can choose either the best direct path from point 1 to exit point 2 (500) or the best indirect path (in this case, via exit point 3, 100+200). This results in a total path metric of $100+200+100=400$. Consider both the exit point's paths and the paths to the exit point in combination. In this case, choose exit point 3, because even though the path to internetwork2 from exit point 3 is 200, the path to exit point 3 from point 1 is only 100, resulting in a total path metric of $100+200=300$. [0236] Relating to one embodiment, FIG. 3a shows an internetwork1 with a network device can decide to redirect flows directly bound for internetwork 4 to another internetwork with a network device, like internetwork 2, or to an internetwork without a network device, like internetwork 3. [0237] Relating to one embodiment, FIG. 3b shows an internetwork1 can have both flow endpoints inside of it, and can have multiple network devices deployed to obtain measurements, calculate metrics, select paths, and alter routing in the internetwork.)

Regarding claim 26

Referring to claim 26 Lloyd discloses all the limitations of claim 26 is described above. Lloyd also discloses wherein the clusters are non-overlapping. (Pg. 20 and 21 A method of populating a plurality of one or more groups with a plurality of one or more network addresses, the method comprising: selecting a plurality of one or more metrics, the plurality of one or more metrics including one or more of path performance, network cost, network usage, geographical proximity, topological proximity, and statistical similarity; creating the plurality of one or more groups in one or more network devices, network devices including one or more of servers, forwarding devices, and routing devices; populating each of the plurality of one or more groups with a subset of the plurality of one or more network addresses based on at least a classification function, the classification function at least partly determined by at least one of the plurality of one or more metrics; and including one or more network addresses from one or more groups of the plurality of one or more groups in a plurality of one or more routing tables distributed across the internetwork.)

Regarding claim 27

Referring to claim 27 Lloyd discloses all the limitations of claim 27 is described above. Lloyd also discloses wherein generalizing the clusters results in associating multiple groups of network addresses with each of at least some of the clusters. (Pg. 1

Art Unit: 2154

Paragraphs[0007] Assessment and Optimization of network traffic is also subjective: performance characteristics toward a given destination vary based on the location from which measurements are made. Each locale must identify their own local performance characteristics. [0008] As the number of valid destination addresses increase, the feasibility of measuring performance characteristics toward all destinations in a persistent and ongoing manner decreases to nil. [0009] Consolidating the set of all individual destination addresses into sets, or ranges of addresses significantly reduces the number of tests that must be made in a given time interval. However, in many cases, even this optimization is insufficient: the current Internet routing tables include over 100,000 different "routes", or address range identifiers.

Regarding claim 28

Referring to claim 28 Lloyd discloses all the limitations of claim 28 is described above. Lloyd also discloses for each cluster, selecting a representative having a predetermined relationship to the seedpoints of the cluster, and associating the representative with each group of network. addresses associated with the cluster. (Pg.13 [0333] In some embodiments of the invention, a single threshold is used, and computed as a certain percentage of the highest score of the paths. In some embodiments of the invention, the threshold is determined by subtracting a fixed quantity to the highest score of the paths.)

Regarding claim 29

Art Unit: 2154

Referring to claim 29 Lloyd discloses all the limitations of claim 29i which is described above. Lloyd also discloses wherein the predetermined relationship of each representative to the seedpoints of the associated cluster is a predetermined relationship of the representative to centroid of the seedpoints. (Pg. 1 [0016] A method of classifying network addresses into appropriate Groups, based on evaluation of a selection of performance, policy, topology, and related criteria. Having defined said Groups, instantiate them in a server or other networking device, such that performance scores for observed traffic can be consolidated within the Group definition.)

Regarding claim 30

Referring to claim 30 Lloyd discloses all the limitations of claim 30 which is described above. Lloyd also discloses wherein the predetermined relationship of each representative to the seedpoints of the associated cluster comprises lying along a network path to a selected one of the seedpoints. (Pg. 1 [0016] A method of classifying network addresses into appropriate Groups, based on evaluation of a selection of performance, policy, topology, and related criteria. Having defined said Groups, instantiate them in a server or other networking device, such that performance scores for observed traffic can be consolidated within the Group definition.) (Pg. 20 and 21 A method of populating a plurality of one or more groups with a plurality of one or more network addresses, the method comprising: selecting a plurality of one or more metrics, the plurality of one or more metrics including one or more of path performance, network cost, network usage, geographical proximity, topological proximity, and statistical similarity; creating the

Art Unit: 2154

plurality of one or more groups in one or more network devices, network devices including one or more of servers, forwarding devices, and routing devices; populating each of the plurality of one or more groups with a subset of the plurality of one or more network addresses based on at least a classification function, the classification function at least partly determined by at least one of the plurality of one or more metrics; and including one or more network addresses from one or more groups of the plurality of one or more groups in a plurality of one or more routing tables distributed across the internetwork.)

Regarding claim 31

Referring to claim 31 Lloyd discloses all the limitations of claim 31 which is described above. Lloyd also discloses wherein selecting a representative for the seedpoints of each cluster includes discarding candidate representatives that do not respond to messages.

(Pg. 1 [0016] A method of classifying network addresses into appropriate Groups, based on evaluation of a selection of performance, policy, topology, and related criteria. Having defined said Groups, instantiate them in a server or other networking device, such that performance scores for observed traffic can be consolidated within the Group definition.)

(Pg. 20 and 21 A method of populating a plurality of one or more groups with a plurality of one or more network addresses, the method comprising: selecting a plurality of one or more metrics, the plurality of one or more metrics including one or more of path performance, network cost, network usage, geographical proximity, topological proximity, and statistical similarity; creating the plurality of one or more groups in one or

Art Unit: 2154

more network devices, network devices including one or more of servers, forwarding devices, and routing devices; populating each of the plurality of one or more groups with a subset of the plurality of one or more network addresses based on at least a classification function, the classification function at least partly determined by at least one of the plurality of one or more metrics; and including one or more network addresses from one or more groups of the plurality of one or more groups in a plurality of one or more routing tables distributed across the internetwork.)

Regarding claim 32

Referring to claim 32 Lloyd discloses all the limitations of claim 32 which is described above. Lloyd also discloses wherein selecting a representative for the seedpoints of each cluster includes discarding candidate representatives that respond to messages with high variability. (Pg.13 [0333] In some embodiments of the invention, a single threshold is used, and computed as a certain percentage of the highest score of the paths. In some embodiments of the invention, the threshold is determined by subtracting a fixed quantity to the highest score of the paths.)

Regarding claim 33

Referring to claim 33 Lloyd discloses all the limitations of claim 33 which is described above. Lloyd also discloses wherein thresholds are employed in ascertaining higher-than-acceptable variability, each threshold being associated with a corresponding source of

Art Unit: 2154

network traffic. (Pg.13 [0333] In some embodiments of the invention, a single threshold is used, and computed as a certain percentage of the highest score of the paths. In some embodiments of the invention, the threshold is determined by subtracting a fixed quantity to the highest score of the paths.)

Regarding claim 34

Referring to claim 34 Lloyd discloses all the limitations of claim 34 which is described above. Lloyd also discloses wherein the representatives are used by an intelligent route controller to select paths for traffic to the destinations, the intelligent route controller being operative to perform periodic measurements to each of the representatives via different connections of the intelligent route controller, and on the basis of the periodic measurements to the representatives, conditionally modify which if the connections is used for traffic sent to the network destinations. (Pg.12 [0235] Relating to an embodiment of FIG. 2, suppose we have a flow at point 1 trying to reach internetwork2, and the paths have the metrics associated with them as annotated, where lower is better and metrics can be added together. Some possibilities for determining the exit point and path to the exit point include (but are not limited to): choose the exit point first and then determine path to the exit point. In this case, choose exit point 2 as the best exit point because its path to internetwork2 has metric 100 and exit point 3 has metric 200. Given that we have chosen exit point 2, we can choose either the best direct path from point 1 to exit point 2 (500) or the best indirect path (in this case, via exit point 3, 100+200). This results in a total path metric of $100+200+100=400$. Consider both the exit point's paths

Art Unit: 2154

and the paths to the exit point in combination. In this case, choose exit point 3, because even though the path to internetwork2 from exit point 3 is 200, the path to exit point 3 from point 1 is only 100, resulting in a total path metric of $100+200=300$.backslash.

[0236] Relating to one embodiment, FIG. 3a shows an internetwork1 with a network device can decide to redirect flows directly bound for internetwork 4 to another internetwork with a network device, like internetwork 2, or to an internetwork without a network device, like internetwork 3. [0237] Relating to one embodiment, FIG. 3b shows an internetwork1 can have both flow endpoints inside of it, and can have multiple network devices deployed to obtain measurements, calculate metrics, select paths, and alter routing in the internetwork.)

Regarding claim 35

Referring to claim 35 Lloyd discloses all the limitations of claim 35 which is described above. Lloyd also discloses wherein the initial grouping of network addresses is established by a set of address prefixes. (Pg. 1 [0016] A method of classifying network addresses into appropriate Groups, based on evaluation of a selection of performance, policy, topology, and related criteria. Having defined said Groups, instantiate them in a server or other networking device, such that performance scores for observed traffic can be consolidated within the Group definition.)

Regarding claim 36

Art Unit: 2154

Referring to claim 36 Lloyd discloses all limitations of claim 36 which is described above. Lloyd also discloses wherein the address prefixes are also employed to establish closeness in the generalizing step. (Pg. 1 [0016] A method of classifying network addresses into appropriate Groups, based on evaluation of a selection of performance, policy, topology, and related criteria. Having defined said Groups, instantiate them in a server or other networking device, such that performance scores for observed traffic can be consolidated within the Group definition.)

Regarding claim 37

Referring to claim Lloyd discloses all the limitations of 37 which is described above. Lloyd also discloses wherein the address prefixes reside in a routing table. (Pg.2 [0051] In some embodiments, the criteria for flow selection can be based on the monitoring of a routing table.)

Regarding claim 38

Referring to claim 38 Lloyd discloses a method of clustering a plurality of network destinations having addresses spanned by a set of address prefixes, comprises: identifying a plurality of seedpoints from among the network destinations, each seedpoint being an active one of the destinations associated with a corresponding at least one of the address prefixes; topologically clustering the seedpoints into groups of topologically similar seedpoints; performing a measurement from a predetermined locations to a seedpoint

Art Unit: 2154

within each group of seedpoints; clustering the seedpoints into clusters based on the measurements, the clusters being selected in a manner achieving a desired trade-off between the number of clusters and the similarity among the measurements for the seedpoints within each cluster; generalizing the clusters based on the address prefixes, the generalizing including conditionally modifying the set of address modified set of address prefixes is associated with a corresponding single one of the clusters. (Pg. 1

Paragraphs[0007] Assessment and Optimization of network traffic is also subjective: performance characteristics toward a given destination vary based on the location from which measurements are made. Each locale must identify their own local performance characteristics. [0008] As the number of valid destination addresses increase, the feasibility of measuring performance characteristics toward all destinations in a persistent and ongoing manner decreases to nil. [0009] Consolidating the set of all individual destination addresses into sets, or ranges of addresses significantly reduces the number of tests that must be made in a given time interval. However, in many cases, even this optimization is insufficient: the current Internet routing tables include over 100,000 different "routes", or address range identifiers. [0010] The performance aspects of network paths to any given address will often have similar characteristics of the network paths to a nearby address. However, the task of identifying address blocks which are acceptably similar in performance characteristic is subtle: network address ranges are typically not directly representative of the underlying network topology. [0011] In most locations, the set of actively used address ranges is a relatively small percentage of the total address ranges available. This set of "active" routes is dynamic, and while technology does exist that can identify the set of active flows, in general these systems

Art Unit: 2154

are constrained by both significant delays in reporting activity, as well as lack of integration into path selection systems.)

Regarding claim 39

Referring to claim 39 Lloyd discloses all the limitations of claim 39 which is described above. Lloyd also discloses wherein the density of the seedpoints in different address prefixes is varied based upon traffic statistics. (Pg. 1 [0007] Assessment and Optimization of network traffic is also subjective: performance characteristics toward a given destination vary based on the location from which measurements are made. Each locale must identify their own local performance characteristics.)

Regarding claim 40

Referring to claim 40 Lloyd discloses all the limitations of claim 40 which is described above. Lloyd also discloses wherein generalizing the clusters includes associating each seedpoint with the longest one of those address prefixes matching the seedpoint. (Pg. 1 [0007] Assessment and Optimization of network traffic is also subjective: performance characteristics toward a given destination vary based on the location from which measurements are made. Each locale must identify their own local performance characteristics.)

Art Unit: 2154

Regarding claim 41

Referring to claim 41 Lloyd discloses all the limitations of claim 41 which is described above. Lloyd also discloses wherein generalizing the clusters results in associating multiple address prefixes with each of a least some of the clusters. (Pg. 1 [0007] Assessment and Optimization of network traffic is also subjective: performance characteristics toward a given destination vary based on the location from which measurements are made. Each locale must identify their own local performance characteristics.)

Regarding claim 42

Referring to claim 42 Lloyd discloses all the limitations of claim 42 which is described above. Lloyd also wherein conditionally modifying the set of address prefixes comprises recursively splitting each address prefix that matches seedpoints from multiple clusters until each resulting address prefix matches seedpoints from only one cluster. (Pg. 7 paragraph [0162] In the implementation described in this document, low resolution in the metric may be considered acceptable, since this tends to diminish the likelihood of reacting to small changes in performance. The goal of some embodiments is to find several routes, none of which are appreciably sub-optimal, and any one of which can be pulled out of service rapidly (with the others taking over) when sudden performance degradations such as packet drop storms are observed. In those embodiments, paths whose measurements vary by insignificant amounts are treated as equal; in some implementations, this can be achieved with a composite metric of low granularity.)

Regarding claim 43

Referring to claim 43 Lloyd discloses all the limitations of claim 43 which is described above. Lloyd also wherein conditionally modifying the set of address prefixes comprises recursively merging address prefixes having greater granularity than the address prefixes in the set of address prefixes until any further merging would result in associating at least one address prefix with seedpoints of multiple clusters. (Pg, 7 paragraph [0162] In the implementation described in this document, low resolution in the metric may be considered acceptable, since this tends to diminish the likelihood of reacting to small changes in performance. The goal of some embodiments is to find several routes, none of which are appreciably sub-optimal, and any one of which can be pulled out of service rapidly (with the others taking over) when sudden performance degradations such as packet drop storms are observed. In those embodiments, paths whose measurements vary by insignificant amounts are treated as equal; in some implementations, this can be achieved with a composite metric of low granularity.)

Regarding claim 44

Referring to claim 44 Lloyd discloses all the limitations of claim 44 which is described above. Lloyd also comprising for each cluster, selecting a representative having a predetermined relationship to the seedpoints of the cluster, and associated with the cluster in the conditionally modified set of address prefixes. (Pg.1 [0009] Consolidating the set

Art Unit: 2154

of all individual destination addresses into sets, or ranges of addresses significantly reduces the number of tests that must be made in a given time interval. However, in many cases, even this optimization is insufficient: the current Internet routing tables include over 100,000 different "routes", or address range identifiers. [0010] The performance aspects of network paths to any given address will often have similar characteristics of the network paths to a nearby address. However, the task of identifying address blocks which are acceptably similar in performance characteristic is subtle: network address ranges are typically not directly representative of the underlying network topology.)

Regarding claim 45

Referring to claim 45 Lloyd discloses all the limitations of claim 45 which is described above. Lloyd also discloses wherein the predetermined relationship of each representative to the seedpoints of the associated cluster comprises lying along a network path to a selected one of the seedpoints. (Pg.1 [0009] Consolidating the set of all individual destination addresses into sets, or ranges of addresses significantly reduces the number of tests that must be made in a given time interval. However, in many cases, even this optimization is insufficient: the current Internet routing tables include over 100,000 different "routes", or address range identifiers. [0010] The performance aspects of network paths to any given address will often have similar characteristics of the network paths to a nearby address. However, the task of identifying address blocks which are acceptably similar in performance characteristic is subtle: network address ranges are typically not directly representative of the underlying network topology.)

Regarding claim 46

Referring to claim 46 Lloyd discloses all the limitations of claim 46 which is described above. Lloyd also discloses wherein the predetermined relationship of each representative to the seedpoints of the associated cluster comprises lying along a network path to a selected one of the seedpoints. (Pg, 7 paragraph [0162] In the implementation described in this document, low resolution in the metric may be considered acceptable, since this tends to diminish the likelihood of reacting to small changes in performance. The goal of some embodiments is to find several routes, none of which are appreciably sub-optimal, and any one of which can be pulled out of service rapidly (with the others taking over) when sudden performance degradations such as packet drop storms are observed. In those embodiments, paths whose measurements vary by insignificant amounts are treated as equal; in some implementations, this can be achieved with a composite metric of low granularity.)

Regarding claim 47

Referring to claim 47 Lloyd discloses all the limitations of claim 47 which is described above. Lloyd also discloses wherein selecting a representative for the seedpoints of each cluster includes discarding candidate representatives that do not respond to messages. (Pg. 8 [0193] In one such implementation, each PoP will be equipped with an implementation of BGP. These BGP sessions should communicate with each other; in some

Art Unit: 2154

implementations, all PoPs can peer in a full mesh of IBGP peerings. In other implementations where the number of PoPs is large so that a full mesh becomes impractical, BGP Route Reflection can be implemented. In such an infrastructure, each PoP can advertise routes to all EC's that it has monitored, giving an address within that PoP as the next hop. Most attributes in that BGP message can be left at default values; in some implementations, only the 32-bit Local Preference field is modified as to communicate the measurements described above for the closeness of the PoP to the EC. Those skilled in the art can identify other fields that can be used for this purpose (e.g., weight, MED, etc.) In implementations that use Local Preference, desirable routes should have higher values; in some embodiments, a scheme based on an ideal number of points from which "demerits" are subtracted, for example when packets are observed to drop, or delay crosses a series of thresholds, can be implemented.

Regarding claim 48

Referring to claim 48 Lloyd discloses all the limitations of claim 48 which is described above. Lloyd also discloses wherein selecting a representative that respond to messages with high variability. (Pg. 8 [0193] In one such implementation, each PoP will be equipped with an implementation of BGP. These BGP sessions should communicate with each other; in some implementations, all PoPs can peer in a full mesh of IBGP peerings. In other implementations where the number of PoPs is large so that a full mesh becomes impractical, BGP Route Reflection can be implemented. In such an infrastructure, each

Art Unit: 2154

PoP can advertise routes to all EC's that it has monitored, giving an address within that PoP as the next hop. Most attributes in that BGP message can be left at default values; in some implementations, only the 32-bit Local Preference field is modified as to communicate the measurements described above for the closeness of the PoP to the EC. Those skilled in the art can identify other fields that can be used for this purpose (e.g., weight, MED, etc.) In implementations that use Local Preference, desirable routes should have higher values; in some embodiments, a scheme based on an ideal number of points from which "demerits" are subtracted, for example when packets are observed to drop, or delay crosses a series of thresholds, can be implemented.

Regarding claim 49

Referring to claim 49 Lloyd discloses all the limitations of claim 49 which is described above. Lloyd also discloses wherein thresholds are employed in ascertaining higher-than-acceptable variability, each threshold being associated with a corresponding source of network traffic. (Pg.13 [0333] In some embodiments of the invention, a single threshold is used, and computed as a certain percentage of the highest score of the paths. In some embodiments of the invention, the threshold is determined by subtracting a fixed quantity to the highest score of the paths.)

Regarding claim 50

Art Unit: 2154

Referring to claim 50 Lloyd discloses all the limitations of claim 50 which is described above. Lloyd also discloses wherein the representatives are used by an intelligent route controller to select paths for traffic to the destinations, the intelligent route controller being operative to perform periodic measurements to each of the representatives via different connections of the intelligent route controller, and on the basis of the periodic measurements to the representatives, conditionally modify which of the connections is used for traffic sent to the network destinations. (Pg.12 [0235] Relating to an embodiment of FIG. 2, suppose we have a flow at point 1 trying to reach internetwork2, and the paths have the metrics associated with them as annotated, where lower is better and metrics can be added together. Some possibilities for determining the exit point and path to the exit point include (but are not limited to): choose the exit point first and then determine path to the exit point. In this case, choose exit point 2 as the best exit point because its path to internetwork2 has metric 100 and exit point 3 has metric 200. Given that we have chosen exit point 2, we can choose either the best direct path from point 1 to exit point 2 (500) or the best indirect path (in this case, via exit point 3, 100+200). This results in a total path metric of $100+200+100=400$. Consider both the exit point's paths and the paths to the exit point in combination. In this case, choose exit point 3, because even though the path to internetwork2 from exit point 3 is 200, the path to exit point 3 from point 1 is only 100, resulting in a total path metric of $100+200=300$.backslash. [0236] Relating to one embodiment, FIG. 3a shows an internetwork1 with a network device can decide to redirect flows directly bound for internetwork 4 to another internetwork with a network device, like internetwork 2, or to an internetwork without a network device, like internetwork 3. [0237] Relating to one embodiment, FIG. 3b shows

Art Unit: 2154

an internetwork1 can have both flow endpoints inside of it, and can have multiple network devices deployed to obtain measurements, calculate metrics, select paths, and alter routing in the internetwork.)

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Ashley d. Turner whose telephone number is 571-270-1603. The examiner can normally be reached on Monday thru Friday 7:30a.m. - 5:00p.m..

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Nathan Flynn can be reached at 571-272-1915. The fax phone number for the organization where this application or proceeding is assigned is 571-270-2603.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

Patent Examiner:

Supervisory Patent Examiner

Ashley Turner

Nathan Flynn

Date: _____

Date: _____

/Nathan J. Flynn/

Supervisory Patent Examiner, Art Unit 2154

